Betting on the FDA: What Option Prices Reveal About Drug Approval Odds

Shanmu Raja, MS*

November 1, 2026

Abstract

When the FDA rules on a drug, the result is close to a coin flip with a known date. The drug is approved or it is not, on a day the agency sets in advance, and the ruling barely moves the rest of the market. That makes it the cleanest forecasting target a public company offers. The options that trade into the decision carry the market's probability that the drug is approved, and this paper recovers that probability from real option prices. We fit a simple model to the option chain: the stock jumps up if the drug is approved and down if it is not, on top of its ordinary day-to-day drift. Fitting that model across strike prices returns the market's approval probability, which we call q. We run it on every FDA decision pending in October 2026. Of 21 decisions, 13 have options liquid enough to fit, and 8 show a clear up-or-down signal; the rest are large firms whose one drug is too small to move the stock, and the model correctly finds no signal there. The eight tell one story. Every one prices the chance of approval below one-half, between 0.08 and 0.36, even though most drugs that reach a decision are approved. That gap is a premium: investors pay to be protected against the bad outcome, so the price-implied probability sits below the true one. These decisions have not happened yet, they resolve between November 2026 and June 2027, so the eight numbers are forecasts recorded in advance that we will score when the outcomes arrive. Three standard pieces of theory sit behind the number: a model-free check on the option term structure confirms the decision is priced at all, the approval probability behaves as a bounded random walk that must settle by the decision date, and a supply-and-demand argument explains why options cost more on the drugs the market fears. A separate set of 24 already-resolved decisions shows the trading side. The option priced a 7.3% move against an actual 3.4%, so options looked expensive, but the part a seller could keep after buying the position back is only 1.3% of the stock, not statistically different from zero ($t = 1.67$), and it rests on a single event. That near-zero result is exactly what the model predicts here, because a liquid options market selects large firms whose single drug does not matter.

1 Introduction

Option prices are forecasts. The price of an option that spans a scheduled event tells you how big a move the market expects, and options struck around a yes-or-no outcome tell you the market's probability of each answer [Breeden and Litzenberger, 1978]. A natural question is whether those forecasts are any good: are they calibrated, and do they beat cheaper forecasts like analyst consensus or prediction markets [Wolfers and Zitzewitz, 2004]? You can only ask that where the

*Draft. Comments welcome. Data and code at the project repository.

outcome is clean and later observed. The FDA decision is the cleanest such target a public company offers, so we ask it there. The tools are standard. One literature measures the gap between the risk-neutral and the real distribution of returns as a risk premium [Carr and Wu, 2009, Pan, 2002, Bollerslev and Todorov, 2011], and another splits option-implied variance into an everyday part and a scheduled-event part, for earnings [Patell and Wolfson, 1979, Dubinsky et al., 2019, Alexiou et al., 2025] and for macro releases [Zhong, 2026]. Our contribution is not the tools but what we point them at: the market's probability of a regulatory verdict, scored against what actually happened.

Four things make an FDA decision unusually clean. The date is fixed ahead of time, because the agency commits to a goal date under the Prescription Drug User Fee Act, so the jump lands on a known day rather than at random. The outcome is close to yes-or-no: an approval or a complete response letter. The direction is unknown in advance, and the two sides are not mirror images, because an approval and a rejection move the stock by different amounts. And the risk is mostly company-specific, since one firm's ruling says little about the market as a whole. Earnings dates share only the first of these, and macro releases share none, which makes the FDA decision the sharpest place to study how options price a one-shot jump of unknown sign on a known date.

We make four contributions. *First*, we recover the option-implied probability of FDA approval, on real data. We model this probability as a random walk that stays between 0 and 1 and must settle on the decision date, borrowing the Jacobi process of Delbaen and Shirakawa [2002] and Ackerer et al. [2018]. We back it out of option prices with no pricing model, by fitting a two-point jump-plus-diffusion mixture, which is just a parametric version of the Breeden–Litzenberger density [Breeden and Litzenberger, 1978, Subramanian, 2004]. Run on every pending FDA decision in October 2026, it gives a clean approval probability for eight of them, and all eight come out below one-half, even though most drugs that reach a decision are approved. *Second*, we show how to explain that probability, and the premium the options charge, with a cross-section on drug fundamentals (Section 5); the prediction, which follows from the fact that informed trading concentrates in small single-drug firms [Bohmann and Patel, 2022, Wu et al., 2024], is that the mispricing is largest in those names. *Third*, we give the theory behind the number: Proposition 1 pulls the event out of the option term structure with a model-free measure, and Proposition 2 explains why options cost more on the drugs that later fail, a crash-insurance premium charged most on the events the market fears, which follows from demand-based option pricing [Gârleanu et al., 2009]. *Fourth*, a pilot of 24 already-resolved decisions shows the trading side: the pre-decision straddle priced a 7.3% move against an actual 3.4%, so it looked rich, but most of that gap is ordinary time value, and the part a seller could keep is 1.3% of the stock, not different from zero across 24 events, which is why the paper is about the forecast and not the trade. Proposition 3 solves the seller's problem, how much to bet, when, and the cost that wipes out the edge, extending Davis et al. [1993] and Leung and Li [2015] to a single fixed-date jump.

How does this differ from existing work? One paper is close enough to address directly. Alexiou et al. [2025] model a scheduled announcement as a jump that makes the option-implied distribution two-humped, and show the straddle seller earns a premium around the event. We do not claim that result: for earnings, options overprice the jump and the seller is paid. Our FDA pilot does not even reproduce it, because the netted premium here is not significant. The contribution is the setting, not a bigger premium. An earnings date comes four times a year, its outcome is a continuous number, and there is no outside probability to check it against. An FDA decision happens once, resolves to yes or no, is recorded, and is separately forecast by analysts and sometimes by prediction markets. That is what makes the *calibration* of the implied probability, not the premium, the question worth asking. We build on the rest of the literature in pieces: Dubinsky et al. [2019] for the variance decomposition, on symmetric earnings jumps; Zhong [2026] for separating scheduled jumps on index

options and macro releases, not single names; Bohmann and Patel [2022] and Wu et al. [2024] for informed trading before FDA events, concentrated in small single-drug firms, which we turn into a cross-sectional prediction; and Jørring et al. [2022] for the economic case that FDA risk is a tradeable binary, though they do not recover it from options. We are not aware of prior work that recovers the FDA approval probability from listed options and scores it.

2 The scheduled-jump model

The model has one idea. Between now and the decision, the stock moves like any other. On the decision day it makes one jump. We do not know the direction in advance: up if the drug is approved, down if it is rejected. We write this down twice. Once for how the stock really behaves, and once for how the option market prices it. Then we compare the two.

The comparison runs through two sets of odds, and the paper is about the gap between them. The *physical* odds, written $\mathbb{P}$, are how things really turn out: how often drugs are approved, and how far the stock really moves. The *risk-neutral* odds, written $\mathbb{Q}$, are the odds baked into option prices. They differ from the real odds because investors pay to avoid risk, so an option can price a bigger move than the stock usually makes. The price of bearing the decision's risk is this gap: what the market charges under $\mathbb{Q}$, minus what the event actually costs under $\mathbb{P}$.

2.1 Setup

Fix a decision at a known date τ. Before that date the stock price S simply diffuses, with an ordinary drift and volatility,

$$\frac{dS_t}{S_{t^-}} = \mu\, dt + \sigma\, dW_t, \qquad t < \tau, \tag{1}$$

and on the date it jumps once, multiplying the price by a random factor that is positive on an approval and negative on a rejection:

$$S_\tau = S_{\tau^-}\,(1+J), \qquad J = \begin{cases} u > 0 & \text{w.p. } p \quad \text{(approval)}, \\ -d,\ d \in (0,1) & \text{w.p. } 1-p \quad \text{(rejection)}. \end{cases} \tag{2}$$

In words: the stock earns its ordinary return until the decision, then takes one discrete step, up by u with probability p or down by d with probability $1-p$. The two-point step is the baseline; Section 5 lets the size vary within each branch. We treat the date τ as known, because the FDA commits to it, and the everyday diffusion as unrelated to the event.

Pricing the option needs the risk-neutral odds, not the real ones. Moving to the risk-neutral measure $\mathbb{Q}$ keeps the same two outcomes but reweights them, replacing the real approval probability p and the real jump sizes with risk-neutral counterparts q and $(u_\mathbb{Q}, d_\mathbb{Q})$. The only discipline is no arbitrage: under $\mathbb{Q}$ the stock must be a fair bet across the jump, which forces

$$\mathbb{E}_\mathbb{Q}[\,1+J \mid \mathcal{F}_{\tau^-}\,] = q\,(1+u_\mathbb{Q}) + (1-q)\,(1-d_\mathbb{Q}) = 1, \tag{3}$$

so the event itself adds no expected drift under the pricing measure.

2.2 What the straddle costs, and the premium we want to measure

A straddle is a call and a put at the same strike. It pays off when the stock moves far in either direction, which is the bet to place when you do not know the direction. The strategy this paper

tests is one trade with it, which we call the wedge straddle. Sell the at-the-money straddle the session before the decision, expiring just after it. Buy it back the session after. What a seller keeps is the move the options priced minus the move that arrived, less the time value that bleeds out over the day held. Take the at-the-money straddle on the session before the ruling, expiring just after it. Its price, as a fraction of the stock, is the move the market is pricing. That price has two parts: the jump the market expects, plus ordinary time value, because the stock keeps moving until the option expires:

$$\frac{\text{Straddle}_{\tau^-}}{S_{\tau^-}} \;=\; \mathbb{E}_{\mathbb{Q}}|J| \;+\; \theta_{\tau^-}, \qquad \theta_{\tau^-} \geq 0. \tag{4}$$

Here $\mathbb{E}_{\mathbb{Q}}|J|$ is the expected size of the jump under the pricing odds, and $\theta_{\tau^-} \geq 0$ is the leftover time value, which disappears only if the option expires the instant the decision lands.

The quantity we are after is the *event premium*: how much more the market prices the jump than the jump actually delivers,

$$\pi \;\equiv\; \mathbb{E}_{\mathbb{Q}}|J| - \mathbb{E}_{\mathbb{P}}|J|, \tag{5}$$

the expected move under the pricing odds minus the expected move that really happens. A naming note, to avoid confusion later. Our premium is about the size of the move, the absolute jump. The "variance risk premium" in the literature is about the square of the move [Carr and Wu, 2009, Bollerslev et al., 2009]. The two are related but not the same, so we reserve "variance risk premium" for the squared version and call ours the event premium.

We cannot read π off directly, but two measurable numbers bracket it. The first is the raw gap between the move the straddle priced and the move that arrived. On average it is

$$\mathbb{E}\left[\frac{\text{Straddle}_{\tau^-}}{S_{\tau^-}} - |J|\right] \;=\; \pi + \mathbb{E}[\theta_{\tau^-}] \;\geq\; \pi, \tag{6}$$

the premium we want plus the leftover time value, so it runs high. A straddle with weeks of life left after the decision still has real value, and subtracting only the realized jump leaves that value in. The second measure strips the time value out. Sell the straddle the session before and buy it back the session after:

$$\mathbb{E}_{\mathbb{P}}\left[\frac{\text{Straddle}_{\tau^-} - \text{Straddle}_{\tau^+}}{S_{\tau^-}}\right]; \tag{7}$$

the buy-back returns the leftover time value. What is left is close to the premium itself, and it is what a seller actually keeps. This is the cleaner number, and the one we headline. Proposition 2 then shows the premium is concentrated in the rejections, and explains why through the price of crash insurance, not larger rejection moves.

2.3 Recovering the approval probability, and a term-structure check

The object the paper is built on is the risk-neutral probability of approval. The full set of option strikes gives it directly, with no pricing model. By Breeden and Litzenberger [1978], the second derivative of the call price in the strike is the risk-neutral density of the stock just after the decision. Because the outcome is binary, that density has two humps: one near the approval price, one near the rejection price. The approval probability q is the mass under the upper hump, and the two humps locate the jump sizes $(u_{\mathbb{Q}}, d_{\mathbb{Q}})$. All three come straight off the density.

In practice, biotech strikes are too coarse for a clean second derivative, so we recover the same object parametrically. We fit a two-point jump-plus-diffusion mixture to the first maturity after the decision. The stock ends at $S_\tau(1+J)$, where J is $u_{\mathbb{Q}}$ with probability q and $-d_{\mathbb{Q}}$ with probability $1-q$, plus a lognormal diffusion of volatility σ. So each option price is q times a Black–Scholes price

struck at the up mode, plus $1-q$ times one struck at the down mode. The martingale condition $q\,u_{\mathbb{Q}} = (1-q)\,d_{\mathbb{Q}}$ removes one parameter. We fit $(q, u_{\mathbb{Q}}, \sigma)$ by inverse-spread-weighted least squares to the out-of-the-money call curve, using puts below spot through put–call parity so the fit rests on the liquid wings. This is just the two-humped density above in parametric form, and it returns q directly. The fit also tests itself. A real binary event makes the mixture beat a single lognormal by a clear margin. A diversified firm whose drug does not matter shows no second hump, and the fit collapses to one mode with q stuck at a boundary. We keep only the decisions where the two-point fit is identified. That is the selection the model expects, because those are the names with large crash exposure.

A coarser check uses only maturities, and confirms an event is priced at all. An option expiring before the decision prices only the everyday diffusion. One expiring just after it prices the diffusion plus the jump. The jump is the step between them. Formally, a model-free variance measure recovers the risk-neutral expected quadratic variation of $\ln S$, which by the log contract of Carr and Madan [1998] and Carr and Wu [2009] is

$$\mathrm{MFIV}(T) \;=\; \underbrace{\sigma^2 T}_{\text{diffusive}} \;-\; \underbrace{2\,\mathbb{E}_{\mathbb{Q}}\big[\ln(1+J)\big]\,\mathbf{1}\{T \geq \tau\}}_{\text{event term}}, \tag{8}$$

using the no-arbitrage condition $\mathbb{E}_{\mathbb{Q}}[J] = 0$ from (3).

Proposition 1 (A term-structure diagnostic for priced event risk)**.** *Suppose the diffusive and jump log-increments are independent under* $\mathbb{Q}$ *and the diffusive variance rate* σ^2 *is continuous across* τ*. Then* $\mathrm{MFIV}(T)$ *has a one-sided jump at the decision date,*

$$V_J \;=\; \lim_{T\downarrow\tau} \mathrm{MFIV}(T) \;-\; \lim_{T\uparrow\tau} \mathrm{MFIV}(T) \;=\; -\,2\,\mathbb{E}_{\mathbb{Q}}\big[\ln(1+J)\big], \tag{9}$$

the event's contribution to expected quadratic variation, which is positive because ln *is concave and* $\mathbb{E}_{\mathbb{Q}}[J] = 0$*. A nonzero* V_J *confirms that the maturities spanning* τ *price an event the shorter ones do not. It is a check that event risk is present, not the recovery of the probability:* V_J *is the log-contract event term, equal to the jump's second moment* $\mathbb{E}_{\mathbb{Q}}[(\ln(1+J))^2]$ *only to leading order and differing at third order in* J*, and the probability comes from the density above, not from this moment.*

Proof. Expected quadratic variation is additive over the diffusive and jump parts, so $\mathrm{MFIV}(T) = \sigma^2 T$ before τ and gains the jump's contribution once its date has passed. The log contract of Carr and Madan [1998], recovered from a strike continuum of out-of-the-money options, prices $2\mathbb{E}_{\mathbb{Q}}[\int (dS/S) - d\ln S]$; across the jump this is $2\mathbb{E}_{\mathbb{Q}}[J - \ln(1+J)]$, which equals $-2\mathbb{E}_{\mathbb{Q}}[\ln(1+J)]$ by $\mathbb{E}_{\mathbb{Q}}[J] = 0$. By continuity of $\sigma^2 T$ the two one-sided limits in (9) differ by exactly that term. □

The two routes reinforce each other. The jump has only two outcomes, so the recovered density has two humps and the short-dated smile is concave near the money [Alexiou et al., 2025]. A concave smile and a term-structure jump in MFIV are two separate signs that the event, not ordinary curvature, is what the options price.

2.4 The downside asymmetry as a price of crash insurance

The raw implied move is larger for events that end in rejection. That pattern is only suggestive, for a reason we give below, but it points to a mechanism. A one-shot two-point model with fixed jump sizes cannot produce a real premium concentrated in rejections. In that model the seller's profit is *smaller* in the branch with the larger move, the opposite of the pattern. The mechanism we propose is instead a selection effect across events, driven by the price of crash insurance. Index

events by i. Let $m_i \equiv \mathbb{E}_{\mathbb{P}}|J_i|$ be the fair expected move and $k_i \geq 0$ the event's crash exposure. The single-name jump cannot be hedged, so by the demand-based pricing of Gârleanu et al. [2009] the dealer who absorbs the demand for protection charges above fair value, by an amount that grows with the unhedgeable risk and the net demand. A risk-averse investor long the sector pays this crash-insurance charge, which we take linear in exposure at a price of crash risk $\lambda \geq 0$, so the risk-neutral straddle is

$$c_i \;\equiv\; \mathbb{E}_{\mathbb{Q}}|J_i| \;=\; m_i + \lambda\, k_i, \tag{10}$$

and events the market fears are more likely to fail, $\mathbb{P}(\text{reject} \mid i) = g(k_i)$ with g nondecreasing. The seller's per-event premium is $\pi_i = c_i - m_i = \lambda k_i$.

Proposition 2 (Downside asymmetry)**.** *Under* (10) *with* $\lambda > 0$ *and* g *nondecreasing and nonconstant:*

(i) the unconditional premium is $\mathbb{E}[\pi_i] = \lambda\,\mathbb{E}[k_i] > 0$*;*

(ii) the harvested premium is larger on rejections, $\mathbb{E}[\pi_i \mid \textit{reject}] \geq \mathbb{E}[\pi_i \mid \textit{approve}]$*, strictly if* g *is strictly increasing on the support of* k*.*

Proof. (i) is the law of iterated expectations. For (ii), since $\mathbb{P}(\text{reject} \mid k) = g(k)$ is nondecreasing, Bayes' rule gives a posterior density of k given rejection whose likelihood ratio to the prior, $g(k)/\mathbb{E}[g]$, is nondecreasing in k; the monotone-likelihood-ratio property implies the rejection-conditional law of k first-order stochastically dominates the approval-conditional law, hence $\mathbb{E}[k \mid \text{reject}] \geq \mathbb{E}[k \mid \text{approve}]$. Multiplying by $\lambda > 0$ gives the claim, since $\pi_i = \lambda k_i$. □

The branch split fits this mechanism but cannot prove it, for a reason worth stating now. Straddles on events that later drew a complete response letter priced a 10.3% move against a realized 4.0%. Those that ended in approval priced a 5.5% move against a realized 3.0%. It is tempting to read the extra charge in the implied move as the crash-insurance premium λk_i of (10). But the raw implied move includes diffusive time value, and feared names carry more diffusive volatility, so their straddles sit higher even with no crash premium at all. In the raw implied move, a pure time-value effect and the crash premium of Proposition 2 look the same. The test that separates them nets the time value out, comparing the short-straddle return by branch. We report it with standard errors in Section 4. At nine rejections it points the way the model predicts, but it is not significant.

2.5 Reading the market's approval probability

The option market also carries a running probability that the drug is approved, and we can track it over time. Call it p_t, the market's approval probability on date t. Two facts pin down how it behaves. It stays between zero and one, because it is a probability. And it has no predictable drift, because if the market could foresee its move, it would have moved already. A process with those two properties is a bounded martingale, and the natural model on $[0, 1]$ is the driftless Jacobi diffusion,

$$dp_t = \nu_t \sqrt{p_t(1 - p_t)}\, dB_t, \qquad p_t \in [0, 1], \tag{11}$$

whose volatility term $\sqrt{p_t(1 - p_t)}$ automatically keeps it inside the interval and whose lack of drift [Delbaen and Shirakawa, 2002, Gourieroux and Jasiak, 2006, Ackerer et al., 2018] matches the no-foresight requirement. One subtlety matters. For p_t actually to settle on a clean yes or no at the

decision date, rather than wandering forever strictly inside the interval, its volatility has to pile up without bound as the date nears,

$$\int_0^\tau \nu_t^2 \, dt \;=\; \infty. \tag{12}$$

A binary that must settle by a fixed date is the case where the implied probability gets jumpier as the date nears, the point Taleb [2018] makes for forecasts that must converge to a verdict. A constant volatility, by contrast, would leave the probability drifting inside the interval and never resolving. The model makes one sharp, testable claim. Because p_t has no predictable drift, the stock should show no predictable abnormal return before the decision, and we test that as a placebo on the pre-ruling window. To read p_t off prices, we use a vertical spread, a pair of options struck around the outcome that together price a one-dollar bet on approval, in the state-price tradition of Breeden and Litzenberger [1978] and the binary-event reading of Subramanian [2004].

2.6 Turning the premium into a trade: how big, how costly, and when

Suppose the premium is real. A seller still faces three questions: how much to bet, how large costs can get before they eat the edge, and when to put the trade on. All three come from one optimization. A seller with standard risk-averse (CRRA) preferences $U(w) = w^{1-\gamma}/(1-\gamma)$ picks an entry time $s < \tau$ and a fraction f of wealth for the short straddle. Its per-event return $\tilde{R}$ has a positive mean $\pi_s = \mathbb{E}_{\mathbb{P}}[\tilde{R}]$ and variance ω^2, net of a round-trip cost κ. Maximizing expected utility,

$$V(t,\cdot) = \sup_{s,\,f} \; \mathbb{E}_{\mathbb{P}}\Big[\, U\big(W_0(1+f\tilde{R})\big) \,\Big|\, \mathcal{F}_t \Big], \tag{13}$$

gives ordinary portfolio choice away from the decision date, with a one-time hit from the jump at it. Its Hamilton–Jacobi–Bellman equation carries a jump term integrating U against the event law (2). Three results fall out.

Proposition 3 (Size, break-even cost, and entry)**.** *(i)* ***Size.*** *The optimal fraction is $f^\star = \pi_s/(\gamma\omega^2) + O(\pi_s^2)$. For log utility with a two-outcome payoff $g_+ > 0 > g_-$ (win on the crush, lose on an adverse gap) occurring with probabilities summing to one, it is exact:*

$$f^\star = \frac{\pi_s}{-\,g_+\,g_-}, \qquad \pi_s = \mathbb{E}_{\mathbb{P}}[\tilde{R}]. \tag{14}$$

(ii) ***Break-even cost.*** *With straddle price $c = \mathbb{E}_{\mathbb{Q}}|J|$ per unit spot, the net expected premium is $\pi_s - \kappa c$, so the trade has positive expected value iff the round-trip spread satisfies*

$$\kappa < \kappa^\star := \pi_s/c. \tag{15}$$

(iii) ***Entry.*** *Let the accrued pre-event P&L over $[s,\tau)$ have physical mean $v\,(\tau - s)$, where $v \geq 0$ is the vega run-up rate, and variance $\beta\,(\tau - s)$, where $\beta > 0$ is the variance accrual rate. The certainty-equivalent gain from holding over the horizon is $\big(v - \frac{\gamma}{2}\beta\big)(\tau - s)$ to leading order, so the optimal entry is the latest possible, $s^\star = \tau^-$, whenever*

$$v < \tfrac{\gamma}{2}\,\beta, \tag{16}$$

and as early as feasible when the inequality reverses, with $v = \frac{\gamma}{2}\beta$ the knife-edge of indifference. The zero-run-up case $v = 0$ falls on the $s^\star = \tau^-$ side.

Proof. (i) The first-order condition $\mathbb{E}_{\mathbb{P}}[\tilde{R}U'(W_0(1+f\tilde{R}))]=0$, expanded to first order in π_s, gives $f^\star=\pi_s/(\gamma\omega^2)$. For log utility and two outcomes, $h(f)=\rho\ln(1+fg_+)+(1-\rho)\ln(1+fg_-)$ has $h'(f)=0$ equivalent to $[\rho g_+ +(1-\rho)g_-]+f\,g_+g_-=0$, i.e. $\pi_s+fg_+g_-=0$, so $f^\star=-\pi_s/(g_+g_-)>0$ since $g_+g_-<0$. (ii) is linear in κ. (iii) A CRRA investor evaluates the added pre-event position by its certainty equivalent, which to second order is mean minus $\frac{\gamma}{2}$ times variance. With mean $v(\tau-s)$ and variance $\beta(\tau-s)$ the certainty equivalent is $\big(v-\frac{\gamma}{2}\beta\big)(\tau-s)$, linear in the horizon $\tau-s$ with slope $v-\frac{\gamma}{2}\beta$. When the slope is negative the holder maximizes by shrinking the horizon to zero, $s^\star=\tau^-$; when positive, by extending it; the sign changes at $v=\frac{\gamma}{2}\beta$. The zero-run-up case has slope $-\frac{\gamma}{2}\beta<0$, giving $s^\star=\tau^-$. □

Part (iii) is why the entry boundary collapses toward the known date τ, unlike the time-homogeneous, Poisson-arrival problems of Merton [1971], Davis et al. [1993], and Leung and Li [2015]. Part (ii) is the decisive test. Even taking the mean short-straddle P&L of $\pi_s\approx 1.3\%$ of spot at face value, against a mean straddle price of $c\approx 7.3\%$, the break-even round-trip spread is $\kappa^\star=\pi_s/c\approx 18\%$ of the straddle. So whether the premium is a real edge or fair pay for jump risk turns on one number: the bid-ask spread a seller actually pays. Single-name biotech straddles routinely quote wider than 18%.

3 Data

We use two datasets. The first is the live set the recovery runs on: every FDA decision still pending at the snapshot, with its option chain. It lets us recover the implied approval probability now and score it when the decisions resolve. The second is a set of resolved decisions with a liquid pre-event straddle, which supplies the realized-move pilot.

3.1 The live pending-decision set

The event universe is the Benzinga FDA calendar, which lists each decision's goal date, drug, sponsor, and ticker. We take every event tagged a PDUFA date whose goal date is after the snapshot of 18 October 2026 and sits on a listed stock, 21 decisions in all. For each, we pull the current option chain from the CBOE delayed-quote feed, which gives bid and ask across strikes and maturities for free, and the spot from Tiingo. We fit the two-point mixture of Section 2.3 to the first maturity after the goal date, building the out-of-the-money call curve from puts and calls through parity. The snapshot is dated, and that is the point. The probabilities are fixed now, and the outcomes land between November 2026 and June 2027, so these are forecasts recorded in advance, not a backfit. One limit is built into the design. Scoring calibration today would need already-resolved decisions, which require a paid options-history feed we do not have. So the scored calibration builds up forward, as the pending set resolves. The free CBOE feed gives current chains, not history.

3.2 The resolved-decision straddle pilot

The pilot sample is every FDA approval and complete response letter on a listed company in the window our resolved options cache covers, late 2024 to September 2026, for which a usable at-the-money straddle traded the session before the decision. We drop decisions with no listed options near the money, or with quotes too stale to trust. What survives is 24 decisions, one per decision,

split 15 approvals and 9 complete response letters. We take one straddle per decision.[1]

For each decision we take the straddle struck nearest the pre-event spot, expiring first after the decision. The implied move is that straddle's price the session before, divided by spot. The realized move is the absolute change in spot from the session before to the session after. The short-straddle P&L is the straddle's price the session before minus its price the session after, again as a fraction of spot. That is what a seller who wrote the straddle into the decision and bought it back the next session would have kept, before costs. The 24 straddles and the script that turns them into every figure below ship with the paper, so the backtest reproduces offline from one file.[2]

One limit governs how every number below should be read. These are option closing prices, not the bid and ask a trader would actually meet. On single-name biotech options the bid-ask gap is wide, and a seller crosses it twice, once to sell and once to buy back. So the premium we report is gross of that spread. Section 4 states the spread it would need to clear, and the conclusion returns to it.

A second feature of the sample governs how far these numbers generalize. Requiring a liquid at-the-money straddle selects firms large enough to carry an options market into the decision. Those are mostly diversified pharmaceutical companies, not single-drug biotechs. The moves give it away: the average realized move on a rejection here is 4.0% of spot, so the drug was immaterial to most of these firms. A single-drug company whose lead asset draws a complete response letter routinely moves twenty to fifty per cent. So the overpricing and the fifteen-of-twenty-four approval rate describe liquid-option, mostly large-cap decisions, and the left tail the sizing discussion warns about is the one this sample leaves out. In the model this is a prediction, not a nuisance. These decisions have crash exposure k_i near zero, so (10) gives a premium $\pi_i = \lambda k_i$ near zero, and the insignificant pilot is exactly what the model forecasts here. The single-drug decisions where k_i is large, and where the premium should show up, are the ones that lack the options liquidity this sample needs. That is the main argument for building the larger universe of Section 5.

4 Results

4.1 The recovered approval probabilities

Table 1 and Figure 1 report the recovery on the live set. Of the 21 pending decisions, 13 have chains liquid enough to fit, and 8 show a clear binary signal. The eight are single-product or lead-asset companies, which is where the model expects a priced event. The names the fit rejects are large diversified sponsors, among them Merck, Bristol Myers Squibb, and BioMarin. Their implied jump is one to four per cent of spot, and q pins at a boundary, because the drug is immaterial and the chain has no second hump to fit. Dropping them is not a judgment call; it is the identification test of Section 2.3 doing its work.

The pattern across the eight is the paper's central finding. Every one prices the chance of approval below one-half, from 0.08 to 0.36, with implied jumps of 16 to 54 per cent of spot. Yet most drugs that reach a decision are approved [Wong et al., 2019]. A risk-neutral probability this far below the base rate is what the structure predicts. By the demand-based argument of Section 2.4, the risk-neutral measure marks the good outcome down and the feared one up, so q sits below the true approval probability by the downside premium. The decisions resolve between November 2026 and June 2027, so this gap is recorded in advance, not fitted after the fact. If these drugs are approved at anything near the historical rate, the distance between that rate and the recovered

[1] An earlier extract of the options cache repeated nine decisions as duplicate rows; treating those repeats as independent observations would understate every standard error, so we deduplicate before any inference.

[2] `data/straddles.csv` and `strategy_backtest.py` in the paper's repository.

q measures the wedge directly. We register the eight here, and will score the Brier score and the reliability diagram as the dates pass.

Ticker	Decision date	$\hat{q}$ (approval)	implied jump	drug
CAPR	2026-11-22	0.27	51%	Deramiocel
COGT	2026-11-30	0.08	16%	bezuclastinib
COGT	2026-12-30	0.24	25%	bezuclastinib
NTLA	2027-03-10	0.26	37%	onvo-z
PHVS	2027-04-23	0.21	41%	deucrictibant
AXSM	2027-05-01	0.30	27%	AXS-12
BBIO	2027-05-08	0.36	30%	encaleret
COGT	2027-06-29	0.28	54%	bezuclastinib

Table 1: Option-implied approval probability $\hat{q}$ recovered from live CBOE chains on 18 October 2026, for the pending PDUFA decisions with an identified binary signal. The implied jump is the risk-neutral expected absolute move $\mathbb{E}_{\mathbb{Q}}|J|$. Every $\hat{q}$ sits below one-half, against a historical base rate in which most drugs reaching a decision are approved. These are pre-registered: outcomes resolve on the listed dates.

4.2 The resolved-decision pilot

Table 2 reports the pilot. Two findings sit at different levels of confidence, and the gap between them is the point. The straddle priced a move about twice the size of the one that arrived, 7.3% of spot against a realized 3.4%. This gross gap is the most stable number in the sample. But it is a statement about the straddle, which still carries time value, not about the decision. A straddle with weeks of life left, against a one-day move, shows a gap like this even when the event premium is zero. The raw implied-move wedge is 3.9% of spot ($t = 3.37$). By (6) it is inflated by the straddle's time value, so it overstates the event premium and is not what a seller keeps. What a seller keeps is the short-straddle return, 1.3% of spot on average. That number does not clear the usual bar: its t-statistic is 1.67, a sign test reads 15 positive events out of 24 ($p = 0.31$), and a bootstrap 95% interval has its lower bound at zero, within 0.01% of spot. The captured premium is positive in point estimate but not distinguishable from zero.

	N	implied move	realized move	implied-move wedge	short-straddle P&L
All events	24	7.3%	3.4%	3.9% (t=3.37)	1.3% (t=1.67)
Approvals	15	5.5%	3.0%	2.4%	0.6%
Rejections	9	10.3%	4.0%	6.3%	2.5%

Table 2: Options and realized moves at FDA decisions, 2024–2026, one observation per decision. Moves and P&L are fractions of spot; the implied-move wedge is the implied move minus the realized move, an upward-biased proxy for the event premium (it contains diffusive time value). All figures are gross of bid-ask spreads.

The result is fragile for a clear reason: it rests on one event. A single December-2024 rejection in AXSM, where the straddle priced a 24% move against a realized 2% and the seller kept 15% of spot, carries much of both figures. Drop it, and the mean short-straddle return falls to 0.7% of spot. The leave-one-out mean ranges only over $[0.7\%, 1.5\%]$ across the other 23 deletions, so

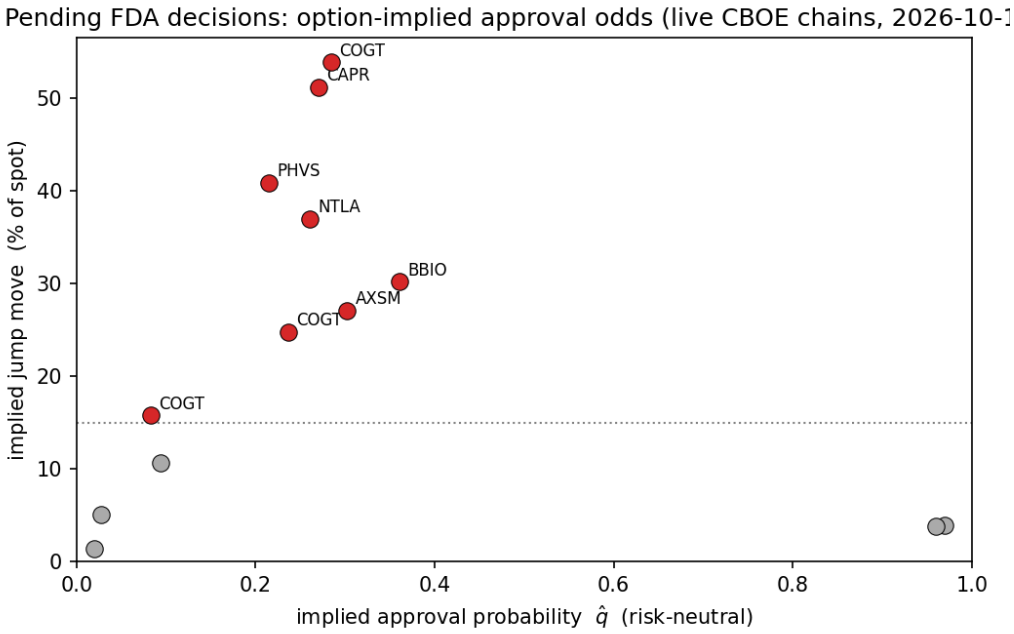


Figure 1: Implied approval probability against implied jump move for the pending decisions. Red points are the identified binary catalysts of Table 1; grey points are names the two-point fit rejects, where the drug is immaterial to a diversified sponsor and $\hat{q}$ is not identified.

no second event rescues it. The robust statistics agree: the median short-straddle return is 0.5% of spot and the 10%-trimmed mean is 0.8%, both well below the 1.3% average. The overpricing wedge holds up better, with a trimmed mean of 3.0% and a median of 1.5%, but even its sign test is only marginal (17 of 24 positive, $p = 0.06$). In short, options overprice the FDA move cleanly, but a seller's captured edge is too small and too concentrated here to call real.

The overpricing is larger for rejections, but this is exactly where the raw wedge misleads. On the nine complete response letters the straddle priced a 10.3% move (standard error 2.1) against a realized 4.0%. On the fifteen approvals it priced 5.5% (standard error 1.0) against a realized 3.0%. That is a raw-wedge gap of 4.8% of spot, with a t near 2.0. On its own it is not evidence of a crash premium, because feared names carry more time value and so a larger raw wedge even with no premium, as Section 2.4 noted. The test that nets the time value out is the short-straddle return by branch: 2.5% on rejections against 0.6% on approvals. With a standard error near 2.0, it is not significant at nine rejections. So the crash-insurance asymmetry fits the sample but is not established by it.

The distribution matters as much as the mean, because selling volatility means a short left tail. Here the best event returned 15.5% of spot (AXSM) and the worst lost 3.3% (an ABBV rejection in April 2026), with a standard deviation of 3.8%. The observed tail is mild. But 24 events over two calm years is a thin basis for it, and the real left tail of a short straddle is a surprise gap of thirty or forty per cent on a small-cap, which this sample never drew. So the observed worst case understates the real one. That is the main risk in reading even the point estimate as an edge.

That gap shows up in the optimal size. Proposition 3(i) applied to these moments gives bets far above one times wealth for log utility, and still above one for a moderately risk-averse seller. These are overbets. They come from a sample variance that leaves out the tail the sizing rule is meant to guard against. The form of the optimum is right, but the number should be set by the worst

plausible gap, and a defined-risk structure that caps the loss is the practical form of the trade.

The break-even cost follows from Proposition 3(ii). Even granting the 1.3% point estimate, with a mean straddle price of 7.3% of spot, the round-trip spread that erases the edge is 18% of the straddle. So whether the premium is real money comes down to one number we do not yet have: the spread a seller would actually pay. Spreads below 18% leave a net premium; spreads above it leave none. On single-name biotech options the quoted spread is often wider still.

Two predictions of the model are not tested by the straddles, and we flag them rather than assert them. The no-drift martingale implies no predictable abnormal return before the decision; the straddle snapshots do not carry the price path, though the stock-side check below gives a coarse version of the test. Proposition 1 needs model-free implied variance across several maturities, and the free feed gives one maturity per event. Both are first on the list for a version built on a fuller options surface.

4.3 A stock-side check on the same wedge

The option is not the only way to ask the question. If the market prices approval below its true odds, a trader who buys the stock the session before the decision and sells it the session after should earn a small positive drift. We test this directly on 94 resolved decisions for which a listed stock price brackets the decision date, taking the close the trading day before and the close the trading day after.[3]

The raw average drift is +1.2% of spot, with a 63% hit rate, but it is not robust, and the reason is the same selection that runs through the paper. A handful of sub-dollar names carry it. A move from one cent to two cents reads as a 45% return and is noise, not a trade a seller could make. Strip the names priced under twenty dollars, the ones that could carry an options market into the decision, and the drift settles to +0.5% of spot, with the mean equal to the median and a t-statistic of 1.5. Figure 2 shows both facts: the large moves live entirely in the cheap, illiquid names, and on the liquid names the drift is small and centered near zero.

That is the answer the straddle gave, now from the stock. On the liquid names where an option trades, neither the option premium nor the stock drift is distinguishable from zero. The tradeable move lives in the small single-drug names, which lack the options liquidity the recovery needs, so the same wall stands whichever instrument we use. The small insignificant drift on liquid names is also what the no-drift martingale predicts, read coarsely through the bracketed move rather than the full path.

5 From recovery to calibration: scoring, scaling, and the cross-section

The recovery is now in hand. Section 4 fixes an implied approval probability for eight pending decisions, and the pilot shows the realized side. What remains is to score those probabilities and to widen the sample. We are explicit about what is done and what is pending. The eight probabilities are registered but not yet scored, because their outcomes resolve over the coming months. A historical backfill that would score calibration today needs the paid options-history feed noted in the data section, which the free current-chain feed does not supply. The calibration battery, the scaling, and the cross-section below are the program that turns the registered forecasts into a measured wedge.

[3] Daily closes from Tiingo; `collect_moves.py` and `data/decision_moves_combined.json` in the paper's repository. The sample skews to approvals, because the calendar under-reports rejections, which if anything flatters a long trade.

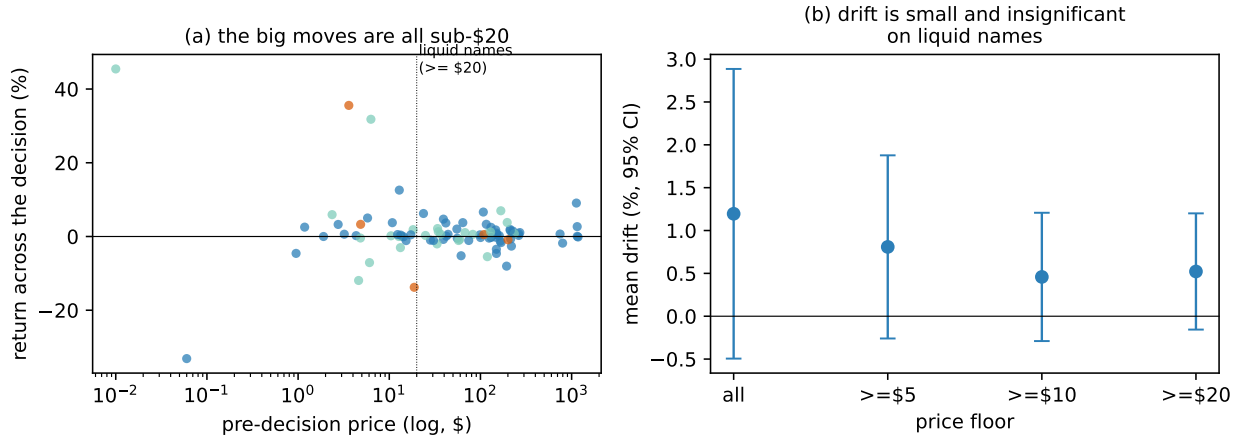


Figure 2: The stock-side wedge on 94 resolved decisions. (a) Return across the decision against the pre-decision price, on a log scale. Every large move sits in the cheap, illiquid names below twenty dollars; the liquid names cluster near zero. (b) Mean long-into-decision drift with a 95% confidence interval, as the price floor rises. The drift is small and not significant on liquid names, and the raw figure is inflated only by sub-dollar noise.

5.1 Scaling the event universe

The binding constraint is sample size, and it can be relaxed. PDUFA decisions are the cleanest events but also the rarest. Two neighboring event types share their structure and add to the sample. Advisory-committee meetings have a fixed date and a recorded vote. Late-stage topline readouts are pre-announced to a window and resolve to a near-binary success or failure. Pooling the three takes the universe from dozens of events a year into the hundreds, the range a calibration test needs, and it lets the implied probability be scored against three kinds of verdict: a regulatory decision, a committee vote, and a trial result.

5.2 Scoring the probability, and what a deviation means

Take the recovered probability $\hat{q}_i$ for event i and the realized outcome $y_i \in \{0, 1\}$. Four standard tools measure how far the market's probability sits from the real one, and in which direction. The eight decisions of Table 1 are the first registered batch, and the tools below score them as each y_i arrives. A reliability diagram bins events by $\hat{q}$ and plots the realized approval frequency in each bin against the bin's mean $\hat{q}$. The Brier score $\frac{1}{n}\sum_i(\hat{q}_i - y_i)^2$ summarizes accuracy, and its decomposition separates calibration from resolution. A regression of the outcome on the implied probability,

$$y_i = a + b\,\hat{q}_i + \varepsilon_i, \tag{17}$$

reads any systematic gap off (a, b). The key is how to read that gap. It is not a test of whether the market is right, and rejecting the unbiased null $a = 0$, $b = 1$ is not evidence of a market error. The model predicts the gap. With $\lambda > 0$ in Proposition 2, the risk-neutral probability sits below the real one, so the unbiased null should be rejected, and the rejection measures the risk-neutral-to-physical wedge, the risk premium, not mispricing. So we read (a, b) as an estimate of that wedge and its slope. Benchmarks sharpen it. Analyst consensus and, where a contract trades, a prediction-market price [Wolfers and Zitzewitz, 2004] are not risk-neutral prices, and they

carry their own biases. Neither is the real probability itself, but the gap between the option-implied probability and those benchmarks brackets the wedge from another side.

5.3 A cross-section on drug fundamentals

An FDA decision, unlike an earnings date, comes with a rich prior, and the cross-section uses it. For each event we observe the therapeutic-area approval rate [Wong et al., 2019], whether the application is a first cycle or a resubmission, whether the drug has drawn a prior complete response letter, the advisory-committee vote where there was one, and whether the sponsor is a single-drug or a diversified firm. Two kinds of test follow. The first regresses the bias $y_i - \hat{q}_i$ on these traits, to find where the market misprices the odds. The natural guess, from the concentration of informed trading in small single-drug firms [Bohmann and Patel, 2022, Wu et al., 2024], is that the bias and the premium are both largest there. The second regresses the event premium and the rejection-side richness of Section 2.4 on the same traits, to test whether the crash-insurance charge scales with the crash exposure k_i the model says drives it. This turns the demand-based asymmetry from an assumption into a measured relationship.

5.4 Measurement extensions

Three measurement upgrades sharpen all of this without changing its shape. Measuring against live bid and ask quotes, not closes, is the only way to settle whether the 18% break-even spread of Proposition 3 is ever cleared. Letting the jump size vary within each branch, instead of the two-point jump, would let the variance decomposition and the recovered probability be estimated together from the full surface. And splitting the premium into a systematic and an idiosyncratic part would test the claim that FDA risk is nearly idiosyncratic, priced lightly by a diversified investor and heavily by the long-the-sector investor of Section 2.4. Decisions that cluster in one therapeutic area should then be treated as correlated, both for inference and because a competitor's ruling moves the same names.

6 Conclusion

Whether the option market charges more for an FDA decision than the decision delivers is the question the resolved pilot cannot settle. Across 24 events the straddle, time value included, priced a 7.3% move against a realized 3.4%, and the raw gap is larger on the drugs that fail. But that is what a straddle carrying weeks of time value shows against a one-day move, even when the event premium is zero. The structure still accounts for the pattern. The demand-based argument predicts a crash-insurance charge heaviest on the drugs the market fears, and the Breeden–Litzenberger density and the term-structure jump confirm that an event is priced at all. What the pilot cannot do is confirm the premium itself, because the raw rejection gap also follows from time value alone, and the netted test is insignificant at nine rejections. What it shows is modest and not new. It is consistent with the scheduled-jump overpricing Alexiou et al. [2025] document for earnings, now seen in a single-name binary setting. The captured return is 1.3% of spot, not different from zero, and it rests on one event, which is what the model predicts for a sample of diversified firms whose single drug is immaterial.

The contribution is the probability, and we recover it. An FDA decision is the rare corporate event whose outcome is binary, dated, later recorded, and separately forecast by analysts and prediction markets. That makes the option-implied probability of approval a forecast you can score, not just state. We recover it with no pricing model, by fitting the two-point mixture to the

live option chain. The eight pending decisions where the fit is identified all price approval below one-half, from 0.08 to 0.36, against a base rate where most drugs reaching a decision are approved. That gap is the risk-neutral-to-physical wedge the structure predicts, the downside premium read straight off the price. Because the decisions are not yet resolved, it is recorded in advance, not fitted after the fact, and the realized approval rate over the coming months will score it. What remains is to scale the universe across advisory-committee votes and late-stage readouts, to the hundreds of observations a full calibration needs, and to test whether the wedge concentrates, as the informed-trading literature suggests, in single-drug firms. The clean question these data pose, and that a larger sample will settle, is whether the market's probability of approval is calibrated, and whether it beats the cheaper forecasts beside it.

References

Damien Ackerer, Damir Filipović, and Sergio Pulido. The Jacobi stochastic volatility model. *Finance and Stochastics*, 22(3):667–700, 2018.

Lykourgos Alexiou, Amit Goyal, Alexandros Kostakis, and Leonidas Rompolis. Pricing event risk: Evidence from concave implied volatility curves. *Review of Finance*, 29(4):963–1007, 2025.

Marc Bohmann and Vinay Patel. Informed options trading prior to FDA announcements. *Journal of Business Finance & Accounting*, 49(7–8):1211–1236, 2022.

Tim Bollerslev and Viktor Todorov. Tails, fears, and risk premia. *Journal of Finance*, 66(6): 2165–2211, 2011.

Tim Bollerslev, George Tauchen, and Hao Zhou. Expected stock returns and variance risk premia. *Review of Financial Studies*, 22(11):4463–4492, 2009.

Douglas T. Breeden and Robert H. Litzenberger. Prices of state-contingent claims implicit in option prices. *Journal of Business*, 51(4):621–651, 1978.

Peter Carr and Dilip Madan. Towards a theory of volatility trading. *Volatility: New Estimation Techniques for Pricing Derivatives*, pages 417–427, 1998.

Peter Carr and Liuren Wu. Variance risk premiums. *Review of Financial Studies*, 22(3):1311–1341, 2009.

M. H. A. Davis, V. G. Panas, and T. Zariphopoulou. European option pricing with transaction costs. *SIAM Journal on Control and Optimization*, 31(2):470–493, 1993.

Freddy Delbaen and Hiroshi Shirakawa. An interest rate model with upper and lower bounds. *Asia-Pacific Financial Markets*, 9(3–4):191–209, 2002.

Andrew Dubinsky, Michael Johannes, Andreas Kaeck, and Norman J. Seeger. Option pricing of earnings announcement risks. *Review of Financial Studies*, 32(2):646–687, 2019.

Nicolae Gârleanu, Lasse Heje Pedersen, and Allen M. Poteshman. Demand-based option pricing. *Review of Financial Studies*, 22(10):4259–4299, 2009.

Christian Gourieroux and Joann Jasiak. Multivariate Jacobi process with application to smooth transitions. *Journal of Econometrics*, 131(1–2):475–505, 2006.

Adam Jørring, Andrew W. Lo, Tomas J. Philipson, Manita Singh, and Richard T. Thakor. Sharing R&D risk in healthcare via FDA hedges. *Review of Corporate Finance Studies*, 11(4):880–922, 2022.

Tim Leung and Xin Li. Optimal mean reversion trading with transaction costs and stop-loss exit. *International Journal of Theoretical and Applied Finance*, 18(3):1550020, 2015.

Robert C. Merton. Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory*, 3(4):373–413, 1971.

Jun Pan. The jump-risk premia implicit in options: Evidence from an integrated time-series study. *Journal of Financial Economics*, 63(1):3–50, 2002.

James M. Patell and Mark A. Wolfson. Anticipated information releases reflected in call option prices. *Journal of Accounting and Economics*, 1(2):117–140, 1979.

Ajay Subramanian. Option pricing on stocks in mergers and acquisitions. *Journal of Finance*, 59 (2):795–829, 2004.

Nassim Nicholas Taleb. Election predictions as martingales: An arbitrage approach. *Quantitative Finance*, 18(1):1–5, 2018.

Justin Wolfers and Eric Zitzewitz. Prediction markets. *Journal of Economic Perspectives*, 18(2): 107–126, 2004.

Chi Heem Wong, Kien Wei Siah, and Andrew W. Lo. Estimation of clinical trial success rates and related parameters. *Biostatistics*, 20(2):273–286, 2019.

Di Wu, Paul Borochin, and Joseph Golec. Informed options trading before FDA drug advisory committee meetings. *Journal of Corporate Finance*, 84:102531, 2024.

Tenghan Zhong. Non-spanning identification of scheduled event risk in option pricing. arXiv preprint, 2026. arXiv:2606.12872.